HOT AND DENSE QUANTUM CHROMODYNAMICS

Panel coordinators: Steffen A. Bass, Duke University
Frithjof Karsch, Brookhaven National Laboratory

Panel Members: P. De Forcrand (Swiss Federal Institute of Technology Zurich), C. DeTar (University of Utah), A. Dumitru (Baruch College and RIKEN-BNL), S. Hands (Swansea University), K. Jansen (*D*eutsches *E*lektronen-*Sy*nchrotron /Zeuthen), H. Meyer (Massachusetts Institute of Technology), B. Mueller (Duke University), P. Petreczky (Brookhaven National Laboratory), C. Schmidt (University of Bielefeld), R. Soltz (Lawrence Livermore National Laboratory), M. Strickland (Gettysburg College), and D. Teaney (Stony Brook)

INTRODUCTION AND CURRENT STATUS

QCD, the theory of the strong force, is one of the central building blocks of the standard model of particle physics that describes the interaction among all known elementary particles (see sidebar 1 in panel report “Cold Quantum Chromodynamics and Nuclear Forces”). QCD uniquely specifies the interactions between the quarks and gluons, and under ordinary conditions such as they exist in the universe today, quarks and gluons do not appear directly as free particles, but are confined into protons and neutrons. QCD predicts that only under extreme conditions of high temperature or of high density (or both) do the quarks and gluons become the most relevant degrees of freedom that dictate the properties of matter. Temperatures that are approximately ten million times larger than those controlling the thermal processes on the surface of the sun, or densities that are approximately ten times larger than those inside a large nucleus, correspond to such extreme conditions. QCD predicts that in such environments, quarks and gluons will behave almost like free particles. This new form of strongly interacting matter is called the quark-gluon plasma (QGP). It existed for a short time in the early universe just after the Big Bang, when matter was still hot and dense. However, after a few microseconds this matter cooled down sufficiently so the thermal conditions no longer allowed for the existence of free quarks and gluons. At this stage, the strongly interacting matter went through a phase transition, reminiscent of the phase transition that occurs when water vapor condenses into a liquid phase. Only after this transition can ordinary matter made out of protons and neutrons be formed in the cosmos.

Deriving detailed predictions of QCD for the properties of matter at high temperature and density is paramount in shaping the current understanding of nuclear matter in general, as well as for understanding the evolution of the early universe. While the properties of nuclear matter at low temperature and moderate densities are well measured, and those of quark-gluon matter at ultra-high temperatures and densities can be rigorously calculated from QCD, little is known about the intermediate regime where the transition between hadronic matter and the QGP occurs. Exploration of this regime, with the aim of mapping the boundaries between different phases of QCD matter, and determining the properties of QCD matter in this domain, is the goal of a large experimental and theoretical program in the United States and internationally. The progress and future goals of this program are described in detail in the *2007 Nuclear Science Long Range Plan: The Frontiers of Nuclear Science*, a report issued by the DOE/NSF Nuclear Science Advisory Committee (NSAC 2007). Large scale computation plays a pivotal role in achieving the goals outlined in this plan. Numerical calculations are required to describe strongly interacting matter in a regime where collective many-particle effects play a dominant role, which helps bridge the gap between analytic calculations performed in the well-defined thermodynamic limit, and environments that are produced in the laboratory.

Experiments studying QCD matter in the domain of interest are currently performed at the RHIC located at BNL, and numerical calculations are performed on leadership-class computers at national laboratories, particularly at the BlueGene/L and Blue-Gene/P computers at Lawrence Livermore National Laboratory (LLNL) and BNL. New experiments are planned for the next decade at RHIC, at the European Organization for Nuclear Research (CERN) at the Large Hadron Collider (LHC), and at the new European Facility for Antiproton and Ion Research (FAIR) to study the properties of matter not only at high temperatures, but also at a high net baryon number density. The primary research goal of the entire experimental and theoretical heavy ion physics program in the United States and worldwide is to provide an answer to one of the central questions raised in the 2007 nuclear science long range plan (NSAC 2007).

What are the phases of strongly interacting matter, and what role do they play in the cosmos?

An important component of the answer to this question is the clarification of whether different regions of the phase diagram of QCD are indeed separated by well-defined *phase transition lines* where properties of matter change abruptly (see sidebar, "QCD Phase Diagram"). Current knowledge of the phase diagram of strongly interacting matter is—to a large extent—based on model-dependent calculations (Stephanov 2006), and very few aspects of the phase-diagram are known from QCD calculations. In particular, it is unknown if a true phase transition and a line of first-order phase transitions actually exist at nonvanishing baryon number density. These questions can be addressed through numerical calculations performed within the framework of LQCD, a discretized version of QCD that is formulated on a four-dimensional grid (lattice) (see sidebar 1 in panel report "Cold Quantum Chromodynamics and Nuclear Forces.") Important steps towards a detailed understanding of the phase diagram of strongly interacting matter have been taken. For vanishing baryon number density, scientists have a reasonable understanding of the transition from ordinary hadronic-matter to the QGP (DeTar, 2008). This transition is a cross-over transition; properties of matter change rapidly in a narrow temperature interval, but the transition is not accompanied by any singular behavior in observables. Reaching this level of understanding for the QCD phase diagram at nonvanishing baryon number density (Schmidt, 2008) requires extreme computational resources.

An important challenge in the studies of strongly-interacting nuclear matter and elementary particles at high temperature is to establish contact between the rigorous Lattice QCD calculations, which establish the equlbrium thermodynamics of this matter, and the properties of the strongly-interacting matter created in heavy ion experiments. This requires an understanding of dynamic properties of hot and dense matter; e.g., transport properties (Meyer 2008) and in-medium properties of hadrons (Asakawa et al., 2001;, Detmold et al., 2009), and an intensive, computationally demanding microscopic modeling of the rapidly expanding and cooling matter created in heavy ion experiments (Nonaka et al.,and Bass 2007). A crucial aspect of such calculations is to determine how the matter—which is originally created in a state far from equilibrium—equilibrates sufficiently rapidly so that a thermodynamic description of its properties becomes possible at short times. Reaching an understanding of the equilibration process in a heavy ion collision requires the modeling of nonequilibrium processes, such as plasma turbulence in a three-dimensional relativistic fluid (TechQM 2008). This is a computationally demanding calculation.

The Phases of Quantum Chromodynamics

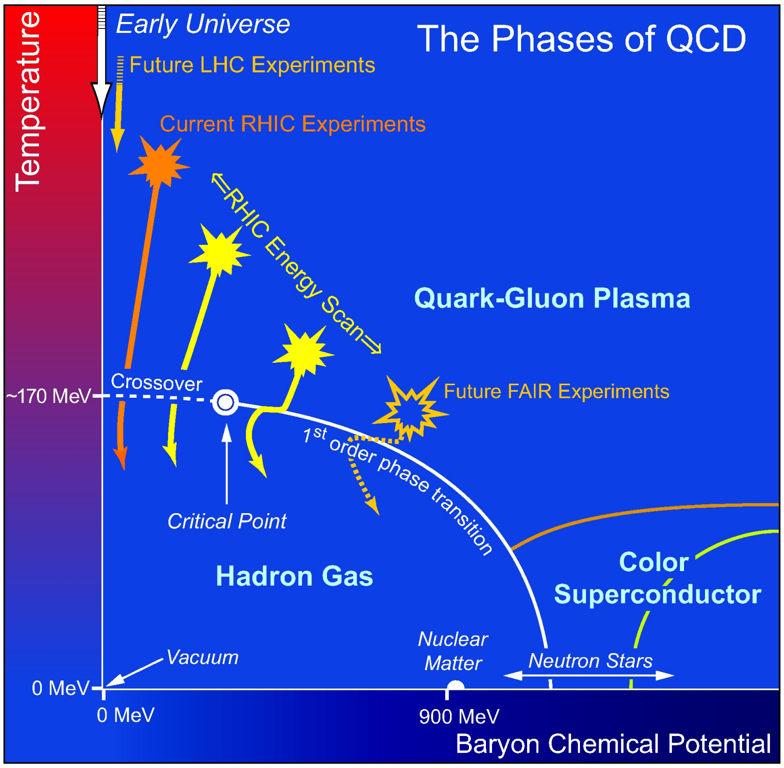

Image from *2007 Nuclear Science Long Range Plan: The Frontiers of Nuclear Science* (NSAC 2007).Exploration of the properties of matter as it existed a few microseconds after the Big Bang in the early universe, and as it may still exist today deep inside neutron stars, is subject to extensive experimental and theoretical investigations. Similar to the phase changes that occur when water is heated or compressed (vapor, fluid, solid), nuclear matter (comprised of protons and neutrons) is expected to undergo drastic changes when heated to extremely high temperatures or compressed to high densities.

Strongly interacting matter at high temperatures and/or densities is expected to consist of “deconfined” quarks and gluons, the elementary building blocks of the theory of strong interactions . While the transition between the low and high temperature regime is not expected to lead to singularities in thermodynamic quantities, this may be different at larger net baryon number densities. A *critical point* is conjectured to mark the density threshold above which the transition between low- and high-temperature regions is accompanied by discontinuities in baryon number densities, and a latent heat is required to dissolve bound states of quarks—the hadrons—into a new form of matter made of free quarks and gluons. At even higher densities, but low temperature, this matter is predicted to exhibit properties akin to those of superconducting materials.

The theoretical studies of hot and dense matter are performed on state-of-the-art supercomputers. Refining these numerical studies into a precision tool capable of establishing the *phase diagram of strongly interacting matter* requires extreme scale computing resources. Extreme scale computing is also required to perform simulations of dynamical processes describing the fascinating experimental results obtained at accelerators such as the RHIC and in future experiments at the European facilities LHC and FAIR.

Creating Hot and Dense Matter in the Laboratory

The only known way to create hot and dense QCD matter under controlled conditions in the laboratory—and to investigate its properties—is to collide two heavy nuclei at velocities close to the speed of light with an accelerator. Since the last decade, accelerators in operation have the capability of creating the temperatures and densities favorable to the formation of a QGP. The overarching goal of these experiments performed, and continue to be performed at these laboratories, is the investigation of the phase diagram of QCD matter, including the deconfined phase, the QGP. The RHIC at Brookhaven National Laboratory (Figure 1) in Long Island, New York, and the accompaniment of detector systems, were built specifically to observe and study the QGP phase of matter.

Figure 1. The RHIC complex at Brookhaven National Laboratory in New York. The complex is comprised of several accelerator facilities joined together to provide beams that are brought into collision in detectors located along the RHIC ring. Image courtesy of Brookhaven National Laboratory (waiting on courtesy permission)

There are four detectors at RHIC: STAR, PHENIX, PHOBOS and BRAHMS. Two are still active, while PHOBOS and BRAHMS completed their operation in 2005 and 2006. Among the two larger detectors, STAR (Figure 2), with its system of time projection chambers covering a large solid angle, is designed for the detection of hadrons; PHENIX is further specialized for detecting rare and electromagnetic interactions.

A typical collision of two gold nuclei, each with momentum of 100 GeV per nucleon, creates a region of QGP matter with a diameter of approximately 10^{-12} cm with a lifetime of approximately 10^{-23} seconds. This QGP fireball then explosively decays into several thousand of particles, which have to be tracked and identified by the detectors (see Figure 3). The particle tracking and characterization of the final state of each collision event poses a significant technological challenge to the RHIC experiments, which can record up to several thousand such events per second. The analysis of these events is used to infer the properties of the transient QGP state.

Creating Hot and Dense Matter in the Laboratory (contd)

Figure 2. The *Solenoid Tracker at RHIC* (STAR) is a detector designed specifically to track the thousands of particles produced by each heavy ion collision at RHIC. Weighing 1200 tons, and as large as a house, STAR is a massive detector. It is used to search for signatures of the QGP, the form of matter that RHIC was designed to create. Image courtesy of Brookhaven National Laboratory (waiting on courtesy permission)

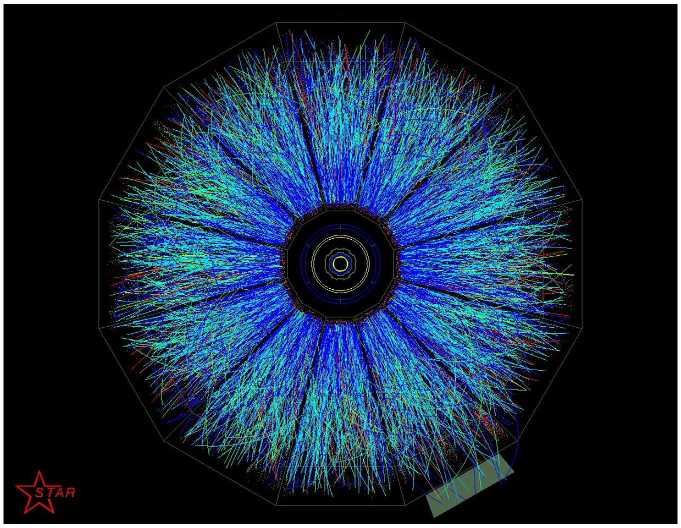

Figure 3. The end view of a collision of two 100-GeV gold beams in the STAR detector at the RHIC at Brookhaven National Laboratory. The beams travel in opposite directions at nearly the speed of light before colliding. Each collision produces thousands of tracks in the detector. Image courtesy of Brookhaven National Laboratory (waiting on courtesy permission)

RHIC began operation in 2000 and is currently the most powerful heavy-ion collider in the world. However, it is expected that the LHC at CERN, completed in 2008, will provide significantly higher energies after it is fully operational in autumn 2009. The planned RHIC-II luminosity upgrade will allow the RHIC and LHC programs to pursue complementary research over the next decade.

The Phase Diagram of Strongly Interacting Matter

The case of vanishing net baryon number density; i.e., the symmetric situation in which the number of particles and antiparticles are identical, plays a special role in the attempt to map out the QCD phase diagram (Stephanov 2006) and to understand phase transitions in strongly interacting matter. Not only does the case of vanishing net baryon number density approximately describe the conditions that existed in the early universe, it is close to conditions that can be studied experimentally in relativistic heavy ion collisions. However, the chiral-limit of QCD, in which the light-quark masses vanish (unphysical values of the quark masses—a theorist's construction) constitutes the only theoretically well understood region of the QCD phase diagram where strongly interacting matter is known to undergo a phase transition. In nature, the quarks are not massless, but two of them (up quarks and down quarks) have a very small mass compared to the scale of chiral-symmetry breaking. It is expected that the thermodynamics of strongly interacting matter reflects many features of this "nearby" chiral phase transition. Today, scientists have reasonably good constraints on the temperature range over which the transition from hadronic-matter to quark-gluon matter occurs. Nonetheless, to provide useful inputs into the modeling of the dense matter created in heavy ion collisions, a reduction is required in the systematic and statistical uncertainties in current determinations of the transition temperature, of the energy density at the transition point, and of the equation of state over the entire temperature range covered by current and future heavy ion experiments.

It is important to firmly establish the existence of a second-order phase transition in the chiral limit of QCD. To date, the expected universal scaling properties of various thermodynamic quantities (Karsch and Laemann et al., 1994) have not been reproduced. This raises concerns about the size of the lattice-spacings used in current LQCD calculations, and could suggest that finer lattices are required to firmly establish the continuum limit of QCD in these calculations. More extensive LQCD calculations, including those with improved discretization schemes for the quarks, are required to improve the current state-of-the-art calculations.

Unlike the regime of vanishing net baryon number density, little is known about the QCD phase diagram at nonvanishing baryon number density through direct numerical calculations. Exploring the structure of the phase diagram at nonvanishing net baryon number density is an outstanding problem that requires numerical approaches quite different from those currently used at zero net baryon number density. As this region of the phase diagram will soon be studied experimentally, it is important to make progress in this area to guide the experimental effort. The computational challenge encountered in this situation is to overcome problems that arise in dealing with extremely high-dimensional integrals that have oscillating integrands. In the context of LQCD calculations at nonzero baryon number density, this is often called the sign problem. Current numerical approaches that try to circumvent this sign problem are promising but require significantly larger computational resources to reach the level of accuracy required to make quantitative statements about the existence, or nonexistence of a critical point, and a line of first-order phase transitions in the QCD phase diagram. At present, calculations using different approaches to circumvent the sign problem lead to conflicting results on the existence of a critical point (Schmidt 2008). These conflicting results may be caused by the drastic approximations that have been introduced in current calculations to make them feasible on today's generation of computers.

Equation of State of Strongly Interacting Matter

The equation of state (DeTar 2008)—or more precisely, the temperature dependence of pressure, energy and entropy density that characterize the static, bulk thermodynamic properties of matter—provides basic information on the relevant degrees of freedom that control properties of strongly interacting matter in different regions of the phase diagram. The rapid rise of energy density over a small temperature interval signals the transition from hadronic to quark-gluon matter. Moreover, the value of the energy density in units of the fourth power of the temperature directly counts the relevant degrees of freedom at a given value of the temperature. Thus, it is easy to determine that the dominant degrees of freedom at low temperature are hadrons, predominantly pions, while at high temperatures the relevant degrees of freedom are the quarks and gluons (see Figure 1). It is more difficult to identify the relevant degrees of freedom in the transition region. Experimental findings at RHIC indicate matter in this regime exhibits properties of a fluid. The existence of various quasi-particle excitations has been postulated to explain the structure of matter in this regime. Also, the remaining significant deviations from the ideal gas behavior of quarks and gluon at high temperatures raise speculation about the QGP properties. Thus, the notion of a quark-gluon liquid has been extensively discussed in the scientific literature, and in view of the experimental findings at RHIC, this liquid has been called a near-perfect liquid. To obtain reliable results on the bulk thermodynamics that allow the verification or falsification of various models of the structure of the high temperature phase of QCD, it is essential to have precise predictions for bulk thermodynamics over a wide temperature range.

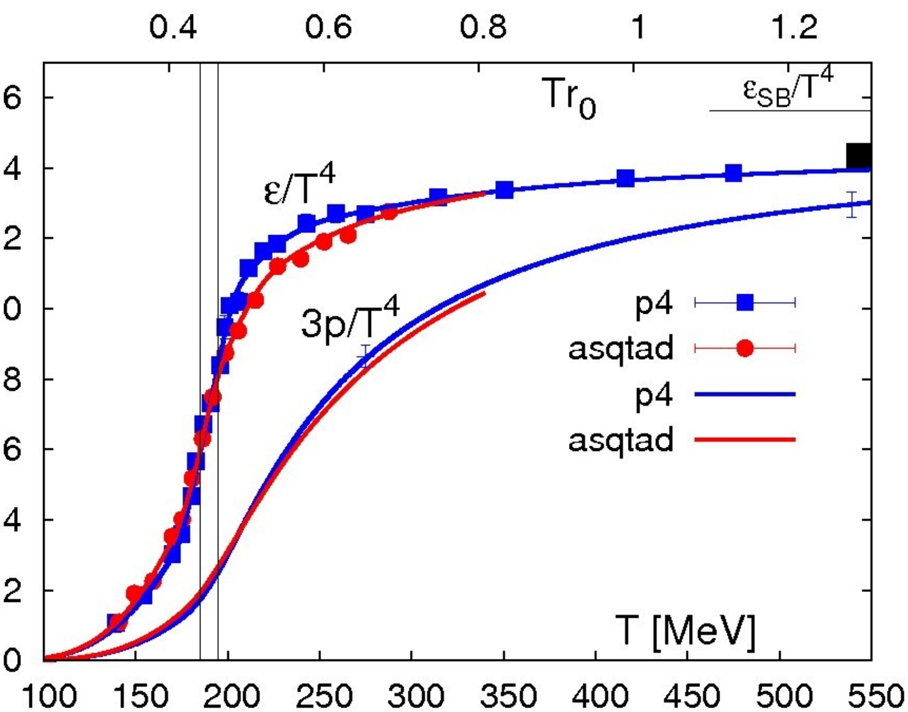

Figure 1. The equation of state of strongly interacting matter calculated with an almost realistic spectrum of quarks for two different discretization schemes (p4 and asqtad). Shown is the energy density and three times the pressure in units of the fourth power of the temperature (Bazavov et al., 2009). Differences between these two quantities indicate the deviation of the equation of state of strongly interacting matter from that of an ideal, noninteracting gas. Source: Bazavov et al. (2009).

The Near-Perfect Liquid Nature of the Quark-Gluon Plasma

In April 2005, Brookhaven National Laboratory announced that scientists at RHIC had created the most "ideal liquid" ever observed in nature. "Elliptic flow" is one of the key features at the center of this discovery. In general, heavy nuclei do not collide head-on, but with their centers displaced by an offset called the "impact parameter" (see Figure 1). This leads to the newly created region of highly compressed hot and dense QCD matter having the cross-sectional shape of an "almond," with pressure gradients pointing outwards perpendicular to its surface. The shape of the compressed zone, in concert with the resulting pressure gradients, leads to the preferential emission of matter along the impact parameter axis of the overlap zone. This phenomenon is termed "elliptic flow," and can be calculated using relativistic fluid dynamics.

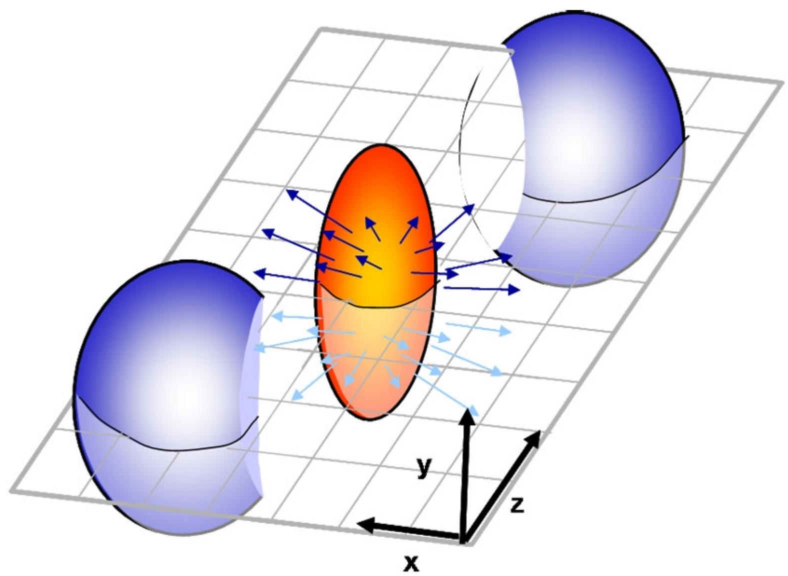

Figure 1. Two nuclei colliding with a nonzero impact parameter create a region of highly compressed QCD matter roughly the shape of an almond. The orientation of the pressure gradients perpendicular to the surface of the almond shape lead to the preferential emission of matter in the xz-plane – this phenomenon is called elliptic flow. Image courtesy of Fermilab (waiting on courtesy permission).

What makes elliptic flow interesting is that it transforms a transient eccentricity of the matter distribution in coordinate space into a measurable eccentricity of matter in momentum space. Calculations using ideal relativistic fluid dynamics have shown that elliptic flow develops early in the collision, at timescales during which the compressed zone was still in the QGP phase. These calculations are in remarkable agreement with RHIC data. The success of ideal relativistic fluid dynamics in describing the elliptic flow observed by scientists conducting the RHIC experiments has led to the conclusion the QGP that is created has the properties of a near-ideal liquid with a very small shear viscosity to entropy-density ratio, η/s. Figure 2 shows a viscous relativistic fluid dynamics calculation of the elliptic flow coefficient, v_2 , as a function of transverse momentum for several values of η/s (Romatschke and Romatschke 2007). Also shown are the experimentally measured values of v_2 and its uncertainties. First exploratory calculations of this ratio in quenched LQCD are consistent with the experimental observations (Meyer 2007). A detailed calculation with light dynamical quarks, however, will require extreme scale computing resources. String theory inspired calculations of this ratio yield a remarkably small lower bound for $\eta/s > 1/4\pi$ in the strong coupling limit of a large class of gauge-theories similar to QCD (Kovtun 2005); unfortunately, such calculations in QCD are not yet possible.

The Near-Perfect Liquid Nature of the Quark-Gluon Plasma (contd)

Most interestingly, however, is that elliptic flow as a manifestation of the near-perfect liquid nature of a system is not restricted to a QGP. The series of pictures in Figure 3 shows elliptic flow occurring in an expanding cloud of ultra-cold lithium atoms being released from an optical trap (O'Hara 2002). Observing the same phenomenon in a system of such different composition and temperature indicates that systems in nature exhibiting elliptic flow may share universal properties.

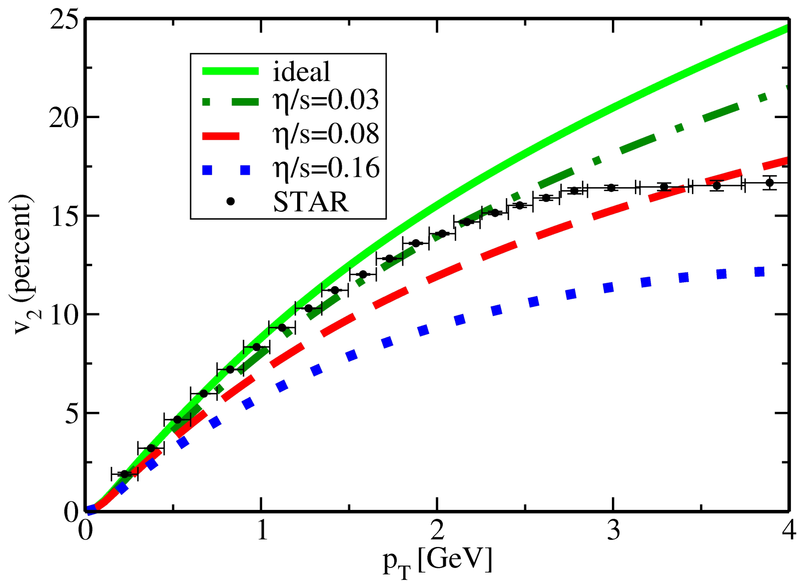

Figure 2. Viscous relativistic fluid dynamics calculations of the elliptic flow coefficient, v2 , as a function of transverse momentum for several different values of the ratio of shear viscosity to entropy-density. The experimental values and associated uncertainties obtained by the STAR collaboration are shown by the black points with error-bars.

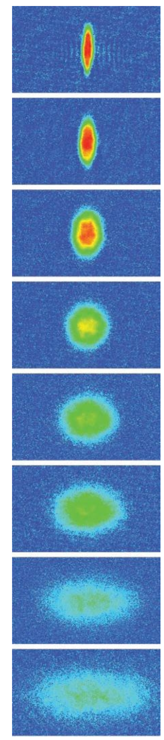

Figure 3. A cloud of ultra-cold fermions (lithium atoms) is released from an optical trap. Due to the initial almond-shape of the trap, the atoms exhibit the same elliptic flow behavior as an expanding QGP formed in ultra-relativistic heavy ion collisions.

To obtain results from LQCD that are also results of QCD, it is necessary to eliminate lattice discretization effects; i.e., perform the continuum limit. To do this in a controlled way requires large scale numerical calculations. As increasing computing resources have become available, the analysis of the QCD equation of state has been refined. The discretization schemes and algorithms used for numerical calculations have become more sophisticated, leading to a reduction of discretization errors but have become computationally more demanding. Today, studies of the equation of state and basic static properties of the QGP are possible with almost realistic parameters. However, these calculations are extremely time consuming. The most advanced study of the equation of state, which just has been completed on leadership-class computers at LLNL and BNL, required approximately 30 teraflop-years to determine the temperature dependence of the energy density and pressure in the limited temperature range currently accessible at RHIC (Bazavov et al. 2009). However, the discretization scheme used in these calculations—improved staggered fermions—is known to violate basic symmetries of QCD at nonzero lattice spacing. Refined calculations at smaller lattice spacings that also cover a larger temperature range, as well as calculations with improved discretization schemes (Jansen 2008)(so-called chiral fermion formulations) are needed in the future to remove the remaining systematic uncertainties in these calculations.

Dynamic Properties of Strongly Interacting Matter

While techniques used to study the static, bulk thermodynamics and the phase diagram of strongly interacting matter are quite advanced, analysis of the dynamic properties remains immature. The temperature dependence of transport coefficients are poorly known, as well as the modification of hadron masses and their widths arising from their interaction with a thermal medium. In-medium modifications of hadron properties are sensitive to the properties of hot and dense matter (Rapp et al.,and Wambach 2000), and can be experimentally studied. Transport coefficients are an important ingredient in the modeling of heavy ion collisions

. It is only recently that a systematic analysis of the experimental data from RHIC using viscous hydrodynamics has been performed (Romatschke et al.,and Romatschke 2007*)*. This analysis has shown that these data are consistent with a small shear-viscosity-to-entropy.

Information on the transport coefficients and the in-medium properties of hadrons are encoded in the spectral functions that characterize the correlation between external sources put into a thermal medium (Hatsuda, 2007). To extract information on spectral functions, high-precision calculations of correlation functions are necessary. A statistical tool, known as the maximum entropy method, can then be used to obtain spectral functions. This approach has many aspects in common with the reconstruction of images from noisy data sets. In the case of QCD, the noisy data are the suitably constructed, numerically evaluated, correlation functions that probe fluctuations of the hot and dense medium. The algorithms used to extract information on transport coefficients and in-medium properties of hadrons from these noisy data are similar to those used in pattern recognition. However, in the context of probing the structure of strongly interacting matter, the main effort is focused on preparing the noisy data. The goal is to reduce the noise level to a point where the filter provided by the maximum entropy method can work efficiently to provide a sharp picture of the structure of the QGP. At present, this approach can only be tested in quenched QCD. Even in this context, the information on correlation functions is barely sufficient to allow for a stable reconstruction of spectral functions. To increase the number of time separations that can give information on the correlation functions without reducing the signal-to-noise ratio in the data, numerical calculations are often performed on anisotropic lattices (lattices that have different lattice spacings in the temporal and spatial directions). In the absence of light quark degrees of freedom, a specific multilevel calculation algorithm is used to reduce the noise level. This approach, however, is not applicable in the presence of dynamical quark degrees of freedom. Much larger computing resources are therefore needed to calculate spectral functions in QCD with all relevant light degrees of freedom.

Microscopic Modeling of Heavy Ion Collisions

Theoretical predictions for thermodynamic properties of strongly interacting matter are based on numerical calculations performed in the framework of LQCD. To explore the properties of hot and dense matter in equilibrium in a realistic system that reflects the properties of QCD correctly, and becomes insensitive to the lattice discretization, requires extraordinary computational resources. as well as the development of new algorithmic concepts. To establish contact between these first principle studies of equilibrium thermodynamics and conditions met in relativistic heavy ion experiments, an additional theoretical interface is required to model the time evolution and cooling of hot and dense matter. The development of a realistic three-dimensional dynamic modeling tool relies on input from studies of equilibrium thermodynamics and, at the same time, is a computationally highly demanding task. An important aspect of this is to obtain a quantitative understanding of the mechanism that leads to the fast equilibration of strongly interacting matter and allows for the observation of thermal effects in heavy ion collisions.

The time evolution of a heavy ion collision at RHIC encompasses several distinct reaction stages, each dominated by very different physical processes (Nonaka and Bass 2007). Figure 2 depicts a schematic view of such a collision: the initial state is comprised of two heavy nuclei (with their wave functions described in terms of elementary quark and gluon degrees of freedom) colliding with each other at approximately 99.9995% speed of light. The system then evolves through a pre-equilibrium stage in which non-Abelian plasma instabilities could drive the system towards equilibrium. Once equilibrated, the now-formed QGP expands hydrodynamically and, in the process, cools down to the critical temperature of QCD, at which point hadronization occurs. The system of newly formed hadrons continues to interact and expand until freeze out, at which point the individual hadrons cease to interact with each other. Each of these reaction stages has its own effective dynamics and computing challenges that need to be addressed to generate a comprehensive understanding of the time evolution of the collision. Hydrodynamic and particle-based Boltzmann codes are well suited to describe the latter three reaction stages—these calculations greatly benefit from the recent advent of grid computing. However, the first two reaction stages pose scientific and computational challenges for which petascale, and ultimately, extreme scale computing will be required.

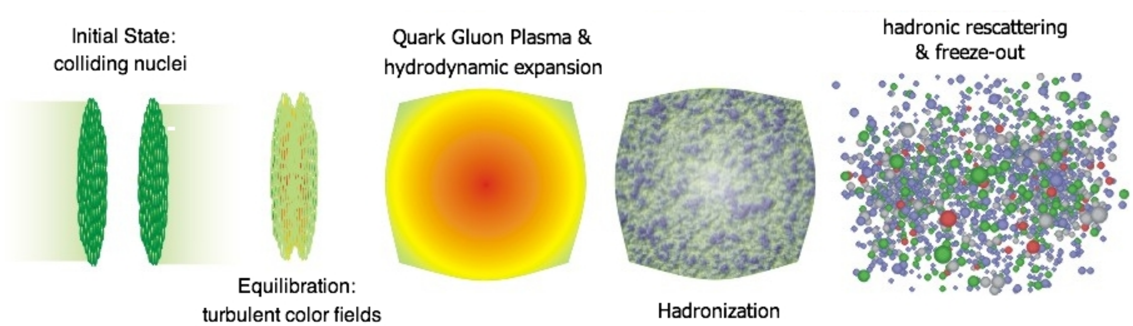

Figure 2. The time evolution of heavy ion collisions at the Relativistic Heavy Ion Collider. The evolution encompasses several distinct reaction stages. Each of these five reaction stages has its own physics and computational challenges that need to be addressed to generate a comprehensive understanding of the time evolution of a collision of relativistic heavy ions. Copyrighted image courtesy of Steffen A. Bass (Duke University).

Precision Calculation of Bulk Thermodynamics

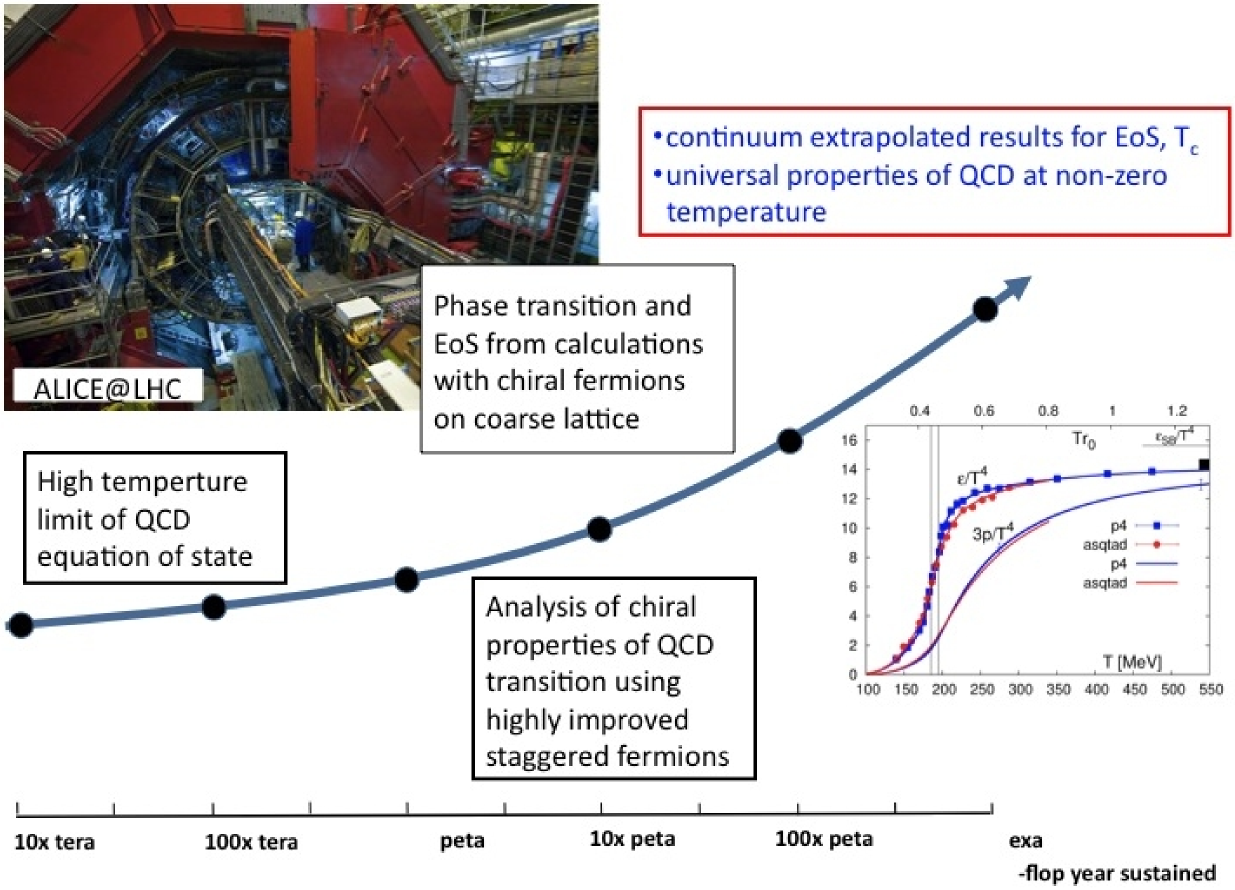

Figure 3. Anticipated highlights for priority research direction, “Precision Calculation of Bulk Thermodynamics” (will correct misspelling in figure)

Basic Science Challenges and Computational Challenges

Establishing the properties of matter in the vicinity of the chiral phase transition, and characterizing their dependences upon the quark masses and the number of quark flavors will provide fundamental insight into the many remarkable features of QCD. It will enable a study of the interplay between the confinement of quarks and gluons and asymptotic freedom, and a study of the role played by chiral symmetry breaking and topological excitations in generating the masses of the hadrons. Furthermore, establishing the properties of strongly interacting matter in the limit of zero net baryon number density is a prerequisite for any further analysis of the QCD phase diagram at nonvanishing baryon number density.

To have complete theoretical control of the thermodynamics of strongly interacting matter in the limit of vanishing baryon number density, it is necessary to extend the existing calculations of the equation of state and basic static properties of hot and dense matter in several respects: 1) extend current knowledge of the equation of state to higher temperatures; 2) establish better theoretical control over the low temperature regime of the equation of state; and 3) better understand the dependence of thermodynamics on the light quark masses to be able to explore the phase transition in the chiral limit.

Equation of State

Basic features of the temperature dependence of the energy density and pressure have already been established through LQCD calculations with rather crude approximations to continuum QCD. Calculations on "coarse" lattices with quark masses that are significantly larger than those of nature have shown that a change in the relevant degrees of freedom occurs over a narrow temperature interval (Karsch et al. 2001). However, even with the most current calculations (Bazavov et al., 2009), full control over the structure of the QCD equation of state has yet to be obtained. At high temperatures, contact has not been established with well-defined analytic calculations. At low temperature, the influence of chiral symmetry breaking and its impact upon the hadronic component of the equation of state has not been established. Moreover, the relevant degrees of freedom that control the structure of the equation of state in the transition region have not been determined. Is the restoration of chiral symmetry of any relevance to the QCD transition, or is the copious production of resonances the driving mechanism that leads to deconfinement and a strongly interacting medium of quarks and gluons at high temperature? To answer these questions, calculations of thermodynamic quantities at higher temperatures must be performed. In addition, lattice discretizations of QCD that respect chiral symmetry or, at least, significantly reduce the influence of its explicit breaking due to the finite lattice spacing, are required in the transition region.

Chiral Fermions

To go beyond the current state-of-the art calculations of the QCD equation of state, it is necessary to use improved discretization schemes for the QCD action that respect all of the symmetries of the continuum theory. These discretization schemes have been developed over several years and continue to be improved through further development. However, these discretization schemes have not been used extensively for numerical studies of QCD to date. This is because they require significantly larger computational resources to perform calculations with sufficiently small statistical uncertainties to allow for a meaningful comparison with the numerical results obtained with nonchiral discretizations.

While highly improved staggered fermion actions like the highly improved staggered quark (HISQ) (Follana et al. 2007) and stout (Morningstar et al.,and Peardon 2004) actions will be used extensively on petaflop computers, truly chiral formulations—such as domain wall and overlap fermion actions (Jansen, 2008) — will require extreme scale computing resources in order for a comprehensive study of chiral aspects of the QCD equation of state. These discretized versions of the QCD action provide significantly better control over the chiral properties of QCD, and thus will be important for analyzing the low temperature and transition region of the static, bulk thermodynamic observables, hadronic screening lengths (Beane 2008a), as well as order parameters that characterize the state of matter at high temperatures. Calculations with chiral fermions will enable the analysis of the universal properties of the transition, such as the scaling behavior of the chiral condensate, its susceptibility as well as quark number susceptibilities, and their fluctuations. This will also allow for them to be related to properties of the equation of state. Further, this will provide a clarification of the relation between the QCD equation of state and the phenomona of deconfinement and chiral symmetry restoration.

High Temperature Limit

Properties of strongly interacting matter at temperatures as large as three to four times the transition temperature will soon be probed experimentally at the LHC at CERN, Switzerland. At these high temperatures, it will be possible to make contact with perturbative calculations in finite temperature and density QCD (Kajantie et al., 2003;, Vuorinen, 2003). This will allow for a cross check between numerical and analytic techniques used in this regime. A reliable numerical calculation of the equation of state and various screening lengths at such high temperatures requires large computational resources because large lattices are needed to control the renormalization of thermodynamic quantities through a proper subtraction of zero temperature observables. This allows for an elimination of otherwise divergent contributions that would prohibit a controlled extrapolation to the continuum limit. Recently developed techniques that minimize the required input from large zero-temperature calculations (Endrodi et al., 2007;, Umeda et al., 2009) have the potential to make these calculations less demanding.

Computational Challenge

Calculations with domain wall fermions or overlap fermions require approximately two orders of magnitude more computational resources than calculations performed with staggered fermions. Prospects for the next generation of studies of bulk thermodynamics based on the staggered fermion discretization scheme have been examined in a white paper written in 2007 by the USQCD collaboration (USQCD 2007). This led to the conclusion that a thorough analysis of the equation of state at temperatures below twice the transition temperature will require approximately 100 sustained teraflop-years. Extending such a study to temperatures twice as high will increase the numerical effort by almost an order of magnitude. A thorough study of the QCD equation of state in the transition from low to high temperature needs to be performed with domain wall or overlap fermions. Such calculations require extreme scale computing resources, as shown in Figure 3.

Outcomes and Impacts

Establishing the properties of strongly interacting matter at vanishing net baryon number density in the chiral limit will define the anchor point for all studies of the QCD phase diagram as a function of temperature and net baryon number density. It will put an indisputable lower bound on the temperature at which hadron matter transforms into a QGP and will establish a reliable starting point for extensions of these calculations into the regime of nonvanishing baryon number density. In combination with calculations using values of light and heavy quark masses as realized in nature, this will quantify the role of chiral symmetry breaking and confinement in the thermodynamics of strongly interacting matter. The equation of state will be the basic equilibrium input to a microscopic description of the rapidly expanding and cooling dense matter formed in a heavy ion collision.

The calculation of the equation of state with physical values for the quark masses will not only have a significant impact on the modeling of heavy ion collisions, it will constrain the range of validity of conventional perturbative calculations at high temperatures and of model building, based on effective theories, at low temperatures.

Quantum Chromodynamics Phase Structure at Non-zero Baryon Density

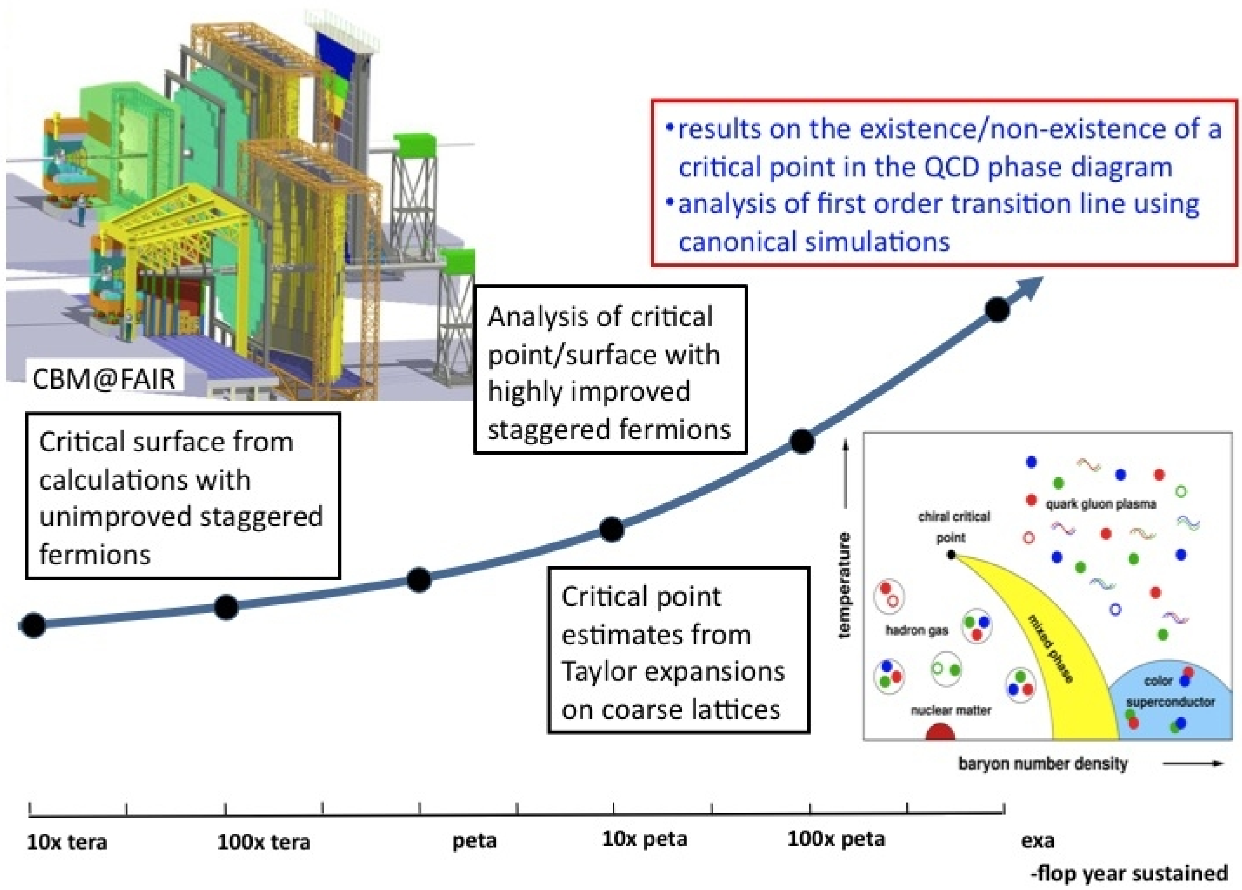

Figure 4. Anticipated highlights for priority research direction, “QCD Phase Structure at Nonzero Baryon Density”

Basic Science Challenges and Computational Challenges

Current studies of the QCD phase diagram and the thermodynamics at nonzero baryon number density are limited to the region of small chemical potential; i.e., small baryon number density. Sensitivity to possible phase transitions at larger values of the chemical potential could arise from conceptually new approaches to the LQCD calculations that overcome the sign problem. This might be achieved through the introduction of auxiliary degrees of freedom that eliminate the oscillating integrands in the QCD partition functions. The complex Langevin approach (Karsch and Wyldet al., 1985;, Aarts and Stamatescu et al., 2008) may eventually lead to such an algorithm that avoids the sign problem. However, to date it has not been successfully implemented in realistic calculations. In the absence of such innovative concepts, currently explored techniques will need to be refined to perform calculations with substantially higher numerical accuracy. These numerical approaches include the Taylor expansion of thermodynamic quantities, such as the pressure and energy density, the analytic continuation of results from numerical calculations performed at imaginary baryon chemical potential, as well as approaches that allow for a projection onto physical states with a fixed baryon number. To use these methods in numerical calculations with physical parameters and improved discretization schemes is challenging and goes beyond currently performed exploratory studies.

Taylor Expansion Techniques

To sufficiently extract information on the existence of phase transitions in the QCD phase diagram from a series expansion of the QCD partition function (Gavai *et al., 2003,* Allton *et al.*, 2003*)*, which directly gives the expansion of the pressure as function of the baryon chemical potential, many expansion coefficients must be determined. This allows for a systematic analysis of the convergence properties of the series and provides insight into the analytic structure of the partition function. The required numerical effort grows rapidly with the order of expansion. Approximately two orders of magnitude increase in computing resources is required to calculate each additional nonvanishing order in the series expansion.

Analytic Continuation

A straightforward way to avoid the sign problem in calculations at nonvanishing baryon number density is to replace the baryon chemical potential with a purely imaginary chemical potential (de Forcrand et al., and Philipsen 2002;, D'Elia et al., and Lombardo 2003). This enables the use of the highly optimized algorithms developed for the calculation of the QCD equation of state at vanishing chemical potential. In particular, it is possible to perform calculations on large lattices with improved actions. However, to extract information on the thermodynamics at nonvanishing baryon number density, extremely precise information is needed on the dependence of thermodynamic observables on the imaginary chemical potential. Only then is it possible to analytically continue the numerical results to the physically relevant finite density regime.

Canonical Ensemble

An attractive, but extremely computationally demanding approach in the numerical studies of strongly interacting matter at nonzero baryon number density, is to perform the calculations directly at a fixed value of the baryon number density (Kratochvila and de Forcrandet al., 2005; Alexandru et al. , 2005). This is in contrast to the approaches discussed above, where calculations are done with an auxiliary control parameter (chemical potential). To perform calculations in the so-called canonical ensemble generally requires the exact calculation of the determinants of large-sparse matrices that is straightforward but computationally demanding. Such calculations may profit from improved eigenvalue solvers optimized for QCD applications.

Color Superconducting Phases

At low temperatures, but large baryon number density, QCD is predicted to become a color superconductor (Rajagopal *et al*., 2000, Alford *et al*., 2008). There may exist several distinct phases, with competing patterns of quark flavor-color-spin-momentum pairings. The existence of such phases may have consequences for understanding the evolution of the early universe and the formation of compact stellar objects.

Very little is known from numerical calculations about the phase structure of strongly interacting matter in this regime (away from the extreme asymptotic limits). First principles calculations in this regime are presently performed only in QCD-like models (Hands, 2007). A direct study within QCD will require the development of new techniques that can manage or circumvent the sign problem. Extreme scale computing resources are required to explore such phases.

Computational Challenge

At present, calculations of Taylor expansions up to the third order in the squared baryon chemical potential require about 100 teraflop-years. Extending these expansions to the fifth order will require resources of 1 exaflop-year. To pursue calculations at these high orders, it is necessary to improve the numerical techniques used to calculate Taylor expansion coefficients. Improved techniques for the inversion of large, sparse matrices (deflation) and the optimization of random source vectors (dilution) are currently being tested and are expected to significantly expedite these calculations. The computational challenges that must be addressed in calculations with imaginary chemical potentials are similar. Quantitative studies of finite density QCD, and a decisive calculation that verifies or excludes the existence of a critical point in the QCD phase diagram, require extreme scale computing resources as shown in Figure 4.***Outcomes and Impact***

Calculations at nonvanishing baryon number density will greatly advance current knowledge of the phase diagram of strongly interacting matter. High-precision calculations of high-order Taylor expansions, as well as accurate calculations with imaginary chemical potential, will provide information on the analytic structure of the QCD partition function. This may allow definitive statements about the density and temperature dependence of the thermodynamics of dense matter to be made, and eventually may determine the location (or rule out its existence) of a critical point in the QCD phase diagram.

These calculations will have an enormous impact on current understanding of properties of strongly interacting matter. These calculations will provide strong constraints on the development of theoretical models for the high-density regime of strongly interacting matter, and will influence the accelerator-based experimental research program in this area.

Transport Coefficients of Quantum Chrmodynamics and Spectral Functions of Hadrons in Medium

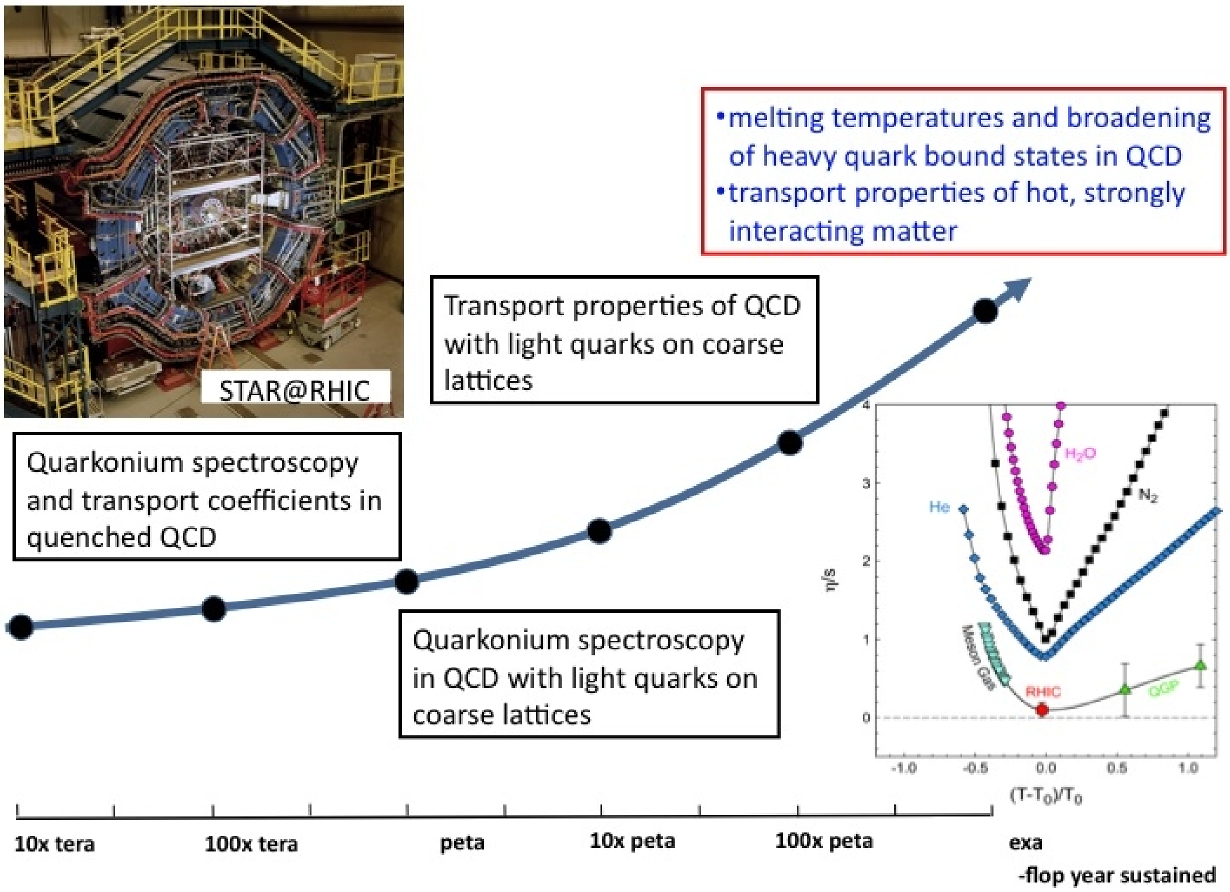

Figure 5. Anticipated highlights for priority research direction, "Transport coefficients of QCD and spectral functions of hadrons in medium." The point labeled RHIC in the figure on the right is a theoretical "estimate."

Basic Science Challenges and Computational Challenges

Numerical calculations of the dynamic properties; i.e., the spectrum of excitations in hot and dense, strongly interacting matter, as well as transport properties of the medium, are presently performed at an exploratory level. To go beyond qualitative statements and to reach a point where quantitative predictions of dynamic properties become feasible, calculations on thermal lattices with unusually large spatial volumes must be performed.

Transport Coefficients

The calculation of transport coefficients, such as the shear and bulk viscosity, that characterize the response of the medium-to-small deviations from its equilibrium state, are particularly difficult. Their calculations formally require taking the zero momentum limit in an infinite spatial volume, which of course, is not possible in numerical calculations. To obtain information on the excitation spectrum of a thermal medium requires accurate calculations of correlation functions at a large set of time separations. The extraction of the (continuous) spectral-function from a finite set of data points is ill-posed (Karsch and Wyld *et al.* 1987). To constrain the class of spectral functions that is consistent with these data, the noise level of the data set at the largest time separations has to be below the ~1% level. Furthermore, the correlation functions have to be calculated at a large number of time separations between sources, making use of correlations between different members of the data set. Thus far, such calculations have only been pursued in quenched QCD (Nakamura and Sakai *et al.*, 2005;, Meyer, 2007). Even in the quenched case, the lattices that were used were too small to obtain reliable results. A petaflop-year of computing resources will be required to complete the studies of transport properties in quenched, and a computation with light dynamical quark degrees of freedom requires extreme scale computing.

In-medium Hadron Masses

The degree of difficulty in calculating the hadronic excitations of the medium that provide information on the in-medium modification of light and heavy quark bound states, is similar to that of the transport coefficients. Hadrons in a thermal bath interact with the medium, and these interactions can lead to the destruction of bound states, and thus the disappearance of the corresponding resonance peaks in the spectral function. Such an effect has been advocated as an experimental signature for the formation of a hot and dense medium in heavy ion collisions (Matsui and Satz *et al*., 1986). Indeed, LQCD calculations of spectral functions at high temperature clearly demonstrate the disappearance of resonance peaks from the hadronic spectral functions (Nakahara et al. 1999). However, to follow the disappearance of these states in hot and dense matter in detail, and locate the melting temperature for various hadronic excitations, requires considerably more computing resources than are currently available. Prior to the disappearance of a state, interactions with the thermal medium will lead to temperature and density dependent shifts of the resonance peaks, as well as a broadening of these peaks. To resolve the structure of spectral functions to such a degree that shifts in resonance peaks and broadening of the spectral curve become statistically significant, accurate numerical results for hadron correlation functions are required. As in the case of calculations of transport coefficients, large lattices are needed to generate information on the correlation functions at many different time separations.

Computational Challenge

The major computational challenge in studies of the excitation spectrum of hot and dense matter is the quest for statistically accurate data on correlation functions on large lattices. These lattices are typically a factor 50 larger than those used in calculations of static, bulk thermodynamics. The size of data samples needed to reach sufficiently small uncertainties in the correlation functions is approximately an order of magnitude larger. Fortunately, such calculations would only be performed at a few selected values of the temperature rather than at the large set of temperature values needed to control properties of the equation of state. Still, this presents a computational challenge, which requires a few petaflop-years to perform calculations within the quenched approximation to QCD. A fully dynamical LQCD calculation, which includes the light quark contributions, will require extreme scale computing resources.

Outcomes and Impacts

Calculations of transport coefficients will provide fundamental insight into the structure of hot and dense matter. It will allow us to quantify aspects of the extent to which the phenomenologically successful modeling of heavy ion collisions has a solid foundation in QCD; i.e., whether a near-equilibrium QGP described by QCD indeed equilibrates rapidly and can be characterized as an almost-perfect fluid. Detailed information on the spectral function would confirm whether or not the QGP is strongly coupled at RHIC, and by varying the temperature in the LQCD calculations, scientists will learn how much the temperature has to be increased before the plasma becomes weakly coupled. This question will be of central importance in comparing the heavy-ion data obtained at the RHIC and the LHC experiments because the temperature in the latter will be about a factor 1.5 to 2 higher.

These calculations will strongly influence the analysis of experimental data obtained in heavy ion collisions.

Equilibration Challenge: From the Color Glass Condensate to the Quark- Gluon Plasma

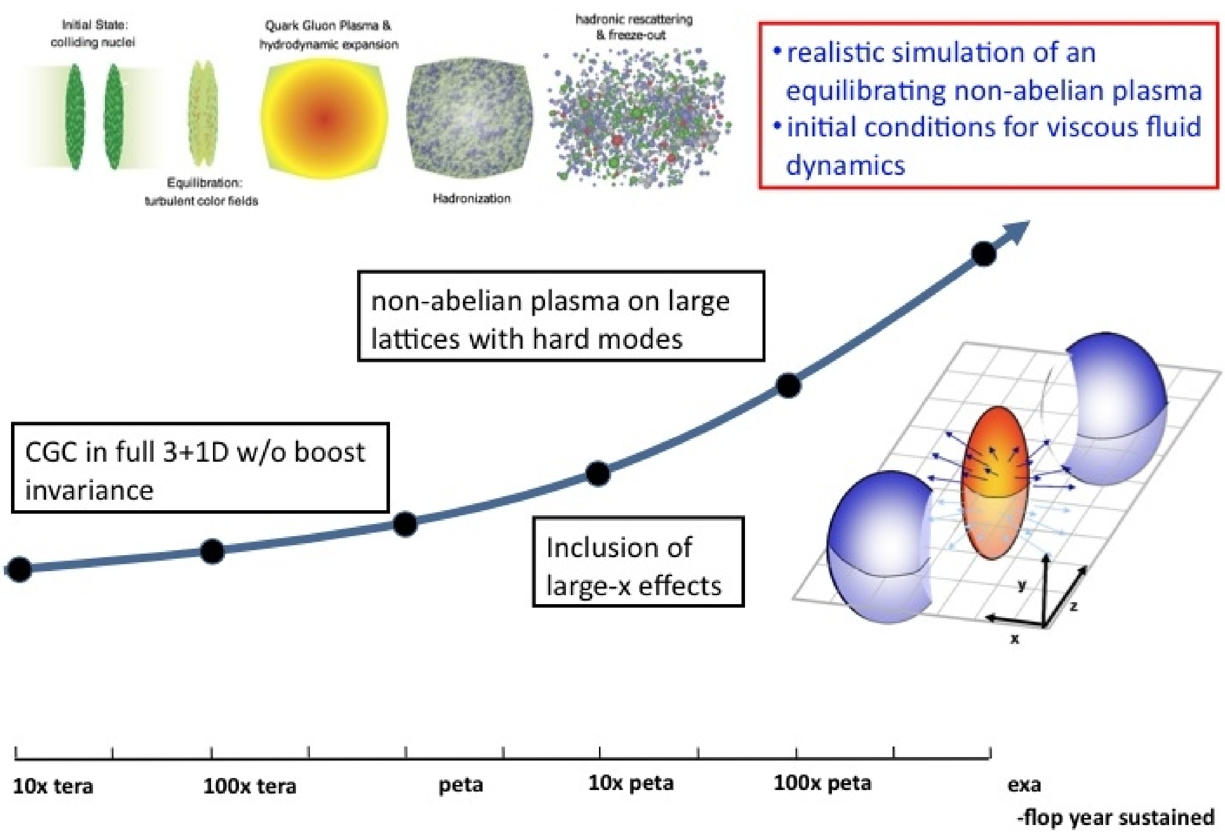

Figure 6. Anticipated research highlights for priority research direction, "Equilibration Challenge: From the Color Glass Condensate to the Quark-Gluon Plasma"***Basic Science Challenges and Computational Challenges***

Determining the time required for a QGP to form after the onset of the collision (i.e., the thermalization time) and determining the physics processes that drive the QGP formation are among the most important outstanding problems in the area of ultrarelativistic heavy ion collisions. The success of ideal hydrodynamics in describing bulk observables—such as the elliptic flow of matter created in noncentral collisions—implies that the matter must have a short thermalization time compared to the overall time scale of the reaction. A proper understanding of QCD dynamics that lead to rapid thermalization is important in determining the transport coefficients and the equation of state of hot QCD.

To describe the approach to equilibrium, the following actions are required: 1) a firm understanding of the initial configuration of partons in the colliding nuclei and the process by which they are liberated from the nucleus at the onset of the collision needs to be acquired; and 2) detailed models and simulations of the processes that occur during the early nonequilibrium phase of the collision leading to the formation of a thermalized QGP need to be developed.

Initial State of the Collision: Color Glass Condensate

A large nucleus moving near the speed of light is a very dense system of gluons. It is believed that nonlinear effects in QCD lead to a saturation of the rapid growth of the gluon density in the colliding nuclei with beam energy and mass number *A* when the phase-space occupation number is (nonperturbatively) large, on the order of $1/\alpha_S$. The effective theory describing this nonlinear regime of QCD is the color glass condensate (CGC) (McLerran 1994a,1994b;, McLerran 1994b, Kovchegov 1996). McLerran and Venugopalan (1994a) proposed an effective action incorporating high-gluon density effects, which amounts to solving the classical Yang-Mills equations where the large-momentum degrees of freedom in the nucleus act as sources of color charge for the small-momentum degrees of freedom.

The classical description of gluon saturation is modified at higher energies due to quantum loop corrections. To this end, a new set of equations, commonly referred to as the JIMWLK (Jalilian-Marian – Iancu – McLerran – Weigert – Leonidov - Kovner) equations (Jalilian-Marian 1997a, Jalilian-Marian 1997b, Iancu 2001), have been derived from a Wilsonian Renormalization Group formalism. They describe the evolution of *n*-point functions in QCD with energy. The resulting equations are an infinite hierarchy of coupled differential equations that are difficult to solve analytically. Nevertheless, they can be written in a form which, in principle, allows them to be solved by lattice gauge-theory techniques (Rummukainen and Weigert 2004).

The CGC has explained several theoretical and phenomenological aspects of high-energy interactions quite successfully. Nevertheless, many important properties of the CGC remain to be addressed quantitatively and tested by comparison with experimental data from RHIC and other colliders.

Thermalization Mechanisms: Plasma Turbulence

In recent years it has been shown that early studies of the driving mechanism for the equilibration of quark-gluon matter in ultra-relativistic heavy ion collisions overlooked a crucial aspect of the dynamics of nonequilibrium plasmas—namely, the possibility of plasma instabilities. Most importantly, it has been shown that these instabilities can produce plasma isotropization and thermalization on a time scale relevant to relativistic heavy ion collisions. The possibility of such nonabelian plasma instabilities was first predicted in the mid-1990s (Mrowczynski 1993) by studying plasmas with an anisotropic momentum-space distribution. The resulting instability has been dubbed the chromo-Weibel instability. In recent years, this theory has received a significant amount of attention because of analytic and numerical advances. The first advance was to show the instabilities predicted by Mrowczynski to be generic and independent of the precise details of the assumed anisotropic distribution function (Romatschke and Strickland 2003). The second advance was the development of a gauge-invariant nonequilibrium effective action that included all nonabelian interactions self-consistently (Mrowczynski et al. 2004). This second advance made it possible to numerically study the dynamics of the chromo-Weibel instability using established real-time lattice gauge theory techniques.

Computational Challenge

The full description of the initial state of a heavy ion collision and subsequent thermalization of the matter requires a code that can self-consistently describe both the earliest periods when the physics of saturation is important, and also the intermediate times when the physics of the chromo-Weibel instability becomes important. To do this requires the real-time solution of the Yang-Mills equation on three-dimensional lattices coupled self-consistently to the Wong equations. Such codes already exist for simplified configurations and expansion scenarios (Bass et al. 1999;, Dumitru et al. 2007;, Schenke et al. 2008). The solution of the full problem requires three-dimensional lattices with a fine lattice-spacing in the longitudinal direction; i.e., solving the classical Yang-Mills equations in real time on a three-dimensional lattice with about 512^3 sites. The field equations describing the small-momentum gluons need to be coupled self-consistently to the Wong equations that describe the propagation of the hard valence sources in the soft background, including energy-momentum conservation. Beyond the classical limit, a simultaneous solution of the rapidity dependence of the JIMWLK measure together with the real-time evolution of the initial fields is required. Beyond that, it is necessary to have a lattice that is capable of describing the dynamics of the chromo-Weibel instability and the subsequent nonabelian cascade to high-momentum modes. The "brute force" way to accomplish this is to ensure the lattice spacing is sufficiently fine. This means using lattices significantly larger than 512^3. As some of the processes (e.g., initial conditions, binary particle collisions, and hard radiation) are stochastic, it will be necessary to average observables over multiple sets of initial conditions (runs). Currently, each run of the simplified calculation takes approximately 1 teraflop-year. Factoring in the higher dimensionality required for the full problem scales this estimate into the tens of petaflop region. Averaging over initial conditions and varying experimental parameters will only be possible with extreme scale computing resources as shown in Figure 6.

Outcomes and Impacts

Extreme-scale computing will deliver real-time calculations of the collision of two heavy ions at high energy, retaining the complete three-dimensional structure of the fields of produced gluons (in impact parameter and rapidity space), energy-momentum conservation, and quantum evolution of the measure. The subsequent real-time evolution of the color fields following the initial impact will clarify the time scales and processes that lead to thermalization and formation of a QGP and the possible role played by non-Abelian gauge-field instabilities (plasma-turbulence). The distribution of the thermalized gluons in the impact parameter and rapidity will provide much-needed initial conditions for hydrodynamic modeling of the late stages of the collision, and could provide information on the equation of state and the viscosity of hot QCD matter. The project would also provide predictions for the effect of early-time nonequilibrium dynamics on important QGP observables such as jet quenching, anomalous transport, and fluctuations.

Applied Mathematics and Computational Issues

Basic problems, which include numerical calculations in QCD at high temperature and density, are common to other fields that make use of the lattice discretized version of the theory of strongly interacting matter. Current programs for these numerical calculations are quite complex, containing several thousand lines of code. They are organized in standardized libraries that have been put in place during the past years with the help of support through the SciDAC software development initiative.

The central, computationally most demanding part in these calculations is related to the inversion of large, sparse matrices, the fermion matrix M (also see discussion in the "Cold Quantum Chromodynamics and Nuclear Forces" panel report). Unlike in other fields, the specific feature of these sparse matrices, which have on the order of a million rows and columns, is that their nonzero entries fluctuate significantly during the calculation of new field configurations. A specific feature of calculations performed in the analysis of dense QCD is this fluctuating background remains frozen for many inversions and only the source vector used to start the inversion process (e.g., based on a conjugate gradient or a similar algorithm) is varying. The need to invert these matrices several hundred times underscores the urgency for improved inversion algorithms that can make efficient use of information collected in the previous inversions. Deflation techniques, domain decomposition, and multigrid methods are being tested and implemented for this purpose. These techniques need to be further developed, and in particular, require optimization for new computing architectures that will be used for extreme scale computing applications.

Enhanced Synergism with Other Subfields of Nuclear Physics and the Broader Physics Community

Establishing reliable quantitative answers to the physics questions addressed in studies of strongly interacting hot and dense matter is important for the understanding of heavy ion collisions and the basic thermodynamic properties of matter. These calculations are fundamentally important for other subfields of nuclear matter, particle physics, astrophysics, and cosmology. Establishing the structure of the QCD phase diagram and the equation of state of strongly interacting matter will have direct consequences for the modeling of the expansion of the early universe. Future large scale astrophysics experiments that aim at a detection of primordial gravitational waves may be sensitive to a direct observation of the phase transition of strongly interacting matter. A possible first-order phase transition at high baryon number density will influence the modeling of compact stars, and may also lead to observable consequences in the cooling pattern of these stars.

Studies of the phase structure of QCD that are necessary to understand hot and dense matter have many features in common with studies performed in statistical physics and material science. In the past, this led to an engaging exchange of ideas on the level of algorithm development, as well as the development of observables and statistical analysis tools needed for the numerical study of phase transitions in general. Numerical algorithms now used in calculations of strongly interacting matter are based on the Monte Carlo algorithm first applied by Metropolis et al. (1953) in studies of the equation of state of molecules. Current versions rely on the molecular dynamics algorithm first used in chemistry (Anderson 1980) for a similar purpose. Statistical analysis tools like the Ferrenberg-Swendsen algorithm (Ferrenberg et al., 1988) and observables like the Binder-cumulant (Binder, 1981) were first applied in statistical physics and are now used as powerful tools to detect phase transitions in strongly interacting matter.

The notorious sign problem faced in numerical studies of QCD at high density led to the development of several new numerical algorithms that have been successfully applied for the simulation of models that are of relevance in statistical physics and material sciences (Chandrasekharan and Wieseet al., 1999).

Numerical calculations of strongly interacting particles have been performed in the framework of LQCD for more than 30 years. These calculations have always demanded large computational resources which, in turn, led several scientists to the development of specific computing hardware. Others developed close collaborations with computer manufacturers. This led to an engaging exchange of ideas and the awareness of the computational needs in LQCD calculations during the design phase of new generations of computers. As a consequence of this close involvement with leading-edge hardware and software developments, many bright scientists who were trained through participation in LQCD calculations now work in the computer and software industry.

Conclusions

Large scale computations within the framework of LQCD have played, and will continue to play, a pivotal role in the exploration of the different phases of hot and dense, strongly interacting matter. For many years these calculations exploited leading-edge computing resources that led to a steady improvement of algorithms, as well as discretization schemes used for these calculations. Given this long-standing experience in large-scale computing, the basic tools exist to address the more complex problems that require extreme scale computing resources. Nonetheless, the exploitation of computing resources, which are three to four orders of magnitude larger than those used today for the calculations of hot and dense matter with LQCD, will require the further development and optimization of existing software tools. An increase in computing power by orders of magnitude also opens up possibilities for the exploitation of new discretization schemes—chiral fermions—which have not yet been used efficiently in studies of hot and dense matter. Furthermore, it will elevate calculations of dynamic properties of strongly interacting matter, which can currently only be performed on an exploratory level, to a stage where calculations with physical parameters will provide reliable input to the microscopic modeling of heavy ion collisions and interpretation of experimental results.

Application of real-time lattice techniques to the nonequilibrium evolution of turbulent fields for the dynamic description of the time evolution of a heavy ion collision will require extreme scale computing resources. This fairly new field of super-computing applications will require significant development of software tools to make optimal use of large scale computing resources.

LQCD calculations are already performed today in large collaborations. The U.S. scientific community is well organized—it manages its own computing resources and coordinates joint proposals for software development projects, and the access to leadership class computing facilities. This valuable infrastructure needs to be maintained and extended. Handling computing resources that are orders of magnitude larger than those used today will require increased workloads and a high degree of organization to process the large amount of data generated in these computational projects. Additional new collaborations need to be formed in the area of dynamic modeling to optimize access to the leadership class computing facilities.

The need for access to extreme scale computing resources for a reliable, quantitative study of properties of hot and dense matter is without question.

PRIORITY RESEARCH DIRECTIONS

HOT AND DENSE QUANTUM CHROMODYNAMICS

PRECISION CALCULATION OF BULK THERMODYNAMICS

Establishing the properties of matter in the vicinity of the chiral phase transition, and characterizing their dependences upon the quark masses and the number of quark flavors, will provide fundamental insight into the many remarkable features of QCD. It will enable a study of the interplay between the confinement of quarks and gluons and asymptotic freedom, and a study of the role played by chiral symmetry breaking and topological excitations in generating masses of the hadrons. Furthermore, establishing the properties of strongly interacting matter in the limit of zero net baryon number density is also a prerequisite for any further analysis of the QCD phase diagram at a nonvanishing baryon number.

To have complete theoretical control of the thermodynamics of strongly interacting matter in the limit of zero baryon number density, it is necessary to extend the existing calculations of the equation of state and basic static properties of hot and dense matter in several respects: Extend scientists' knowledge of the equation of state to higher temperatures

1. Establish better theoretical control over the low temperature regime of the equation of state

2. Better understand the dependence of thermodynamics on light quark masses to be able to explore the phase transition in the chiral limit.

Basic features of the temperature dependence of the energy density and pressure have already been established through LQCD calculations with rather crude approximations to continuum QCD. Calculations on "coarse" lattices with quark masses that are significantly larger than those in nature have shown that a change in the relevant degrees of freedom occurs over a narrow temperature interval (Karsch *et al.,* 2001). However, even with the most current (Bazavov et al. , 2009), full control over the structure of the QCD equation of state is yet to be obtained. At high temperatures, contact is yet to be established with well defined analytic calculations. At low temperatures, the influence of chiral symmetry breaking, and its impact upon the hadronic component of the equation of state is yet to be established. Moreover, the relevant degrees of freedom that control the structure of the equation of state in the transition region are yet to be determined. Is the restoration of chiral symmetry of any relevance to the QCD transition, or is the copious production of resonances the driving mechanism that leads to deconfinement and a strongly interacting medium of quarks and gluons at high temperature? To answer these questions, calculations of thermodynamic quantities at higher temperatures must be performed. In addition, lattice discretizations of QCD that respect chiral symmetry or, at least, significantly reduce the influence of its explicit breaking due to the finite lattice spacing, are required in the transition region.

To go beyond the current state-of-the art calculations of the QCD equation of state, it is necessary to use improved discretization schemes for the QCD action that respect all of the symmetries of the continuum theory. These discretization schemes have been developed over several years and continue to be improved through further development. However, they have not been used extensively for numerical studies of QCD to date. This is because they require significantly larger computational resources to perform calculations with sufficiently small statistical uncertainties to allow for a meaningful comparison with the numerical results obtained with nonchiral discretizations.

While highly improved staggered fermion actions like the highly improved staggered quark (HISQ) (Follana et al. 2007) and stout (Morningstar and Peardon 2004) actions will be used extensively on petaflop computers, truly chiral formulations—such as domain wall and overlap fermion actions (Jansen, 2008) — will require extreme scale computing resources in order for a comprehensive study of chiral aspects of the QCD equation of state. These discretized versions of the QCD action provide significantly better control over the chiral properties of QCD, and thus will be important for analyzing the low temperature and transition region of the static, bulk thermodynamic observables, hadronic screening lengths (Beane 2008a), as well as order parameters that characterize the state of matter at high temperatures. Calculations with chiral fermions will enable the analysis of the universal properties of the transition, such as the scaling behavior of the chiral condensate, its susceptibility as well as quark number susceptibilities, and their fluctuations. This will also allow for them to be related to properties of the equation of state. Further, this will provide a clarification of the relation between the QCD equation of state and the phenomona of deconfinement and chiral symmetry restoration.

Properties of strongly interacting matter at temperatures as large as three to four times the transition temperature will soon be probed experimentally at the LHC at CERN, Switzerland. At these high temperatures, it will become possible to make contact with perturbative calculations in finite temperature and density QCD (Kajantie et al., 2003;, Vuorinen, 2003). This will allow for a cross check between numerical and analytic techniques used in this regime. A reliable numerical calculation of the equation of state and various screening lengths at such high temperatures requires large computational resources because large lattices are needed to control the renormalization of thermodynamic quantities through a proper subtraction of zero temperature observables. This allows for an elimination of otherwise divergent contributions that would prohibit a controlled extrapolation to the continuum limit. Recently developed techniques that minimize the required input from large zero-temperature calculations (Endrodi et al., 2007;, Umeda et al., 2009) have the potential to make these calculations less demanding.

Establishing the properties of strongly interacting matter at zero net baryon number density in the chiral limit will define the anchor point for all studies of the QCD phase diagram as function of temperature and net baryon number density. It will put an indisputable lower bound on the temperature at which hadron matter transforms into a QGP and will establish a reliable starting point for extensions of these calculations into the regime of nonvanishing baryon number density. In combination with calculations using values of light and heavy quark masses as realized in nature, this will quantify the role of chiral symmetry breaking and confinement in the thermodynamics of strongly interacting matter. The equation of state will be the basic equilibrium input to a microscopic description of rapidly expanding and cooling dense matter formed in a heavy ion collision.

The calculation of the equation of state with physical values for the quark masses will not only have a significant impact on the modeling of heavy ion collisions, it will constrain the validity range of conventional perturbative calculations at high temperatures and of model building, based on effective theories at low temperatures.

QUANTUM CHROMODYNAMICS PHASE STRUCTURE AT NONZERO BARYON DENSITY

Current studies of the QCD phase diagram and the thermodynamics at nonzero baryon number density are limited to the region of small chemical potential; i.e., small baryon number density. Sensitivity to possible phase transitions at larger values of the chemical potential could arise from conceptually new approaches to the LQCD calculations that overcome the sign problem. This might be achieved through the introduction of auxiliary degrees of freedom that eliminate the oscillating integrands in the QCD partition functions. The complex Langevin approach (Karsch et al.,and Wyld 1985;, Aarts and Stamatescuet al., 2008) may eventually lead to such an algorithm that avoids the sign problem. However, as yet it has not been successfully implemented in realistic calculations. In the absence of such innovative concepts, currently explored techniques will need to be refined to perform calculations with substantially higher numerical accuracy. These numerical approaches include the Taylor expansion of thermodynamic quantities, such as the pressure and energy density, the analytic continuation of results from numerical calculations performed at imaginary baryon chemical potential, as well as approaches that allow for a projection onto physical states with a fixed baryon number. To use these methods in numerical calculations with physical parameters and improved discretization schemes is challenging and goes beyond any currently performed exploratory studies.

To extract information on the existence of phase transitions in the QCD phase diagram from a series expansion of the QCD partition function (Gavai and Guptaet al., 2003; Allton et al., 2003), which directly gives the expansion of the pressure as a function of the baryon chemical potential, sufficiently many expansion coefficients must be determined. This enables a systematic analysis of the convergence properties of the series, and provides insight into the analytic structure of the partition function. The required numerical effort grows rapidly with the order of expansion. Approximately a two orders of magnitude increase in computing resources is required to calculate each additional nonvanishing order in the series expansion.

A direct approach to avoid the sign problem in calculations at nonvanishing baryon number density is to replace the baryon chemical potential with a purely imaginary chemical potential (de Forcrand et al.,and Philipsen 2002;, D'Elia et al.,and Lombardo 2003). This allows for the use of the highly optimized algorithms developed for the calculation of the QCD equation of state at vanishing chemical potential. In particular, it is possible to perform calculations on large lattices with improved actions. However, to extract information on the thermodynamics at nonvanishing baryon number density, scientists need extremely precise information on the dependence of thermodynamic observables on the imaginary chemical potential. Only then is it possible to analytically continue the numerical results to the physically relevant finite density regime.

An attractive, but extremely computationally demanding approach in the numerical studies of strongly interacting matter at nonzero baryon number density, is to perform the calculations directly at a fixed value of the baryon number density (Kratochvila and de Forcrandet al., 2005; Alexandru et al., 2005). This is in contrast to the above-discussed approaches where calculations are performed with an auxiliary control parameter (chemical potential). To perform calculations in the so-called canonical ensemble generally requires the exact calculation of the determinants of large sparse matrices, which is straightforward but computationally demanding. Such calculations may profit from improved eigenvalue solvers optimized for QCD applications.

At low temperatures, but large baryon number density, QCD is predicted to become a color superconductor (Rajagopal and Wilczek et al., 2000;, Alford et al. 2008). There may exist several distinct phases, with competing patterns of quark flavor-color-spin-momentum pairings. The existence of such phases may have consequences for understanding the evolution of the early universe and the formation of compact stellar objects.

Little is known from numerical calculations about the phase structure of strongly interacting matter in this regime (away from the extreme asymptotic limits). First-principle calculations in this regime are, at present, performed only in QCD-like models (Hands 2007). A direct study within QCD will require the development of new techniques that can manage or circumvent the sign problem. Extreme scale computing resources are required for an exploration of such phases.

Calculations at nonvanishing baryon number density will greatly advance our knowledge of the phase diagram of strongly interacting matter. High-precision calculations of high-order Taylor expansions, as well as accurate calculations with imaginary chemical potential, will provide information on the analytic structure of the QCD partition function. This may allow definitive statements about the density and temperature dependence of the thermodynamics of dense matter to be made, and eventually may determine the location or rule out the existence of a critical point in the QCD phase diagram.

These calculations will have an enormous impact on scientists' understanding of properties of strongly interacting matter. They will provide strong constraints on the development of theoretical models for the high-density regime of strongly interacting matter and will influence the accelerator-based experimental research program in this area.

TRANSPORT COEFFICIENTS OF QUANTUM CHRMODYNAMICS AND SPECTRAL FUNCTIONS OF HADRONS IN MEDIUM

Numerical calculations of the dynamic properties; i.e., the spectrum of excitations in hot and dense, strongly interacting matter, as well as transport properties of the medium, are still performed at an exploratory level. To go beyond qualitative statements and reach a stage where quantitative predictions of dynamic properties become feasible, calculations on thermal lattices with unusually large spatial volumes have to be performed.

The calculation of transport coefficients, such as the shear and bulk viscosity, that characterize the response of the medium-to-small deviations from its equilibrium state, are particularly difficult. Their calculation formally requires taking the zero momentum limit in an infinite spatial volume, which of course, is not possible in numerical calculations. To obtain information on the excitation spectrum of a thermal medium requires accurate calculations of correlation functions at a large set of time separations. The extraction of the (continuous) spectral-function from a finite set of data points is ill-posed (Karsch and Wyldet al., 1987). To constrain the class of spectral functions that is consistent with these data, the noise level of the data set at the largest time separations has to be below the ~1% level. Furthermore, the correlation functions have to be calculated at a large number of time separations between sources, making use of correlations between different members of the data set. So far, this program has only been pursued in quenched QCD (Nakamura and Sakaiet al., 2005;), Meyer, 2007). Even in the quenched case, the lattices that were used were too small to obtain reliable results. A petaflop-year of computing resources will be required to complete the studies of transport properties in quenched, and a computation with light dynamical quark degrees of freedom requires extreme scale computing.

The degree of difficulty in calculating the hadronic excitations of the medium that provide information on the in-medium modification of light and heavy quark bound states, is similar to that of the transport coefficients. Hadrons in a thermal bath interact with the medium, and these interactions can lead to the destruction of bound states and thus the disappearance of the corresponding resonance peaks in the spectral function. Such an effect has been advocated as an experimental signature for the formation of a hot and dense medium in heavy ion collisions (Matsui and Satzet al., 1986). Indeed, LQCD calculations of spectral functions at high temperature clearly demonstrate the disappearance of resonance peaks from the hadronic spectral functions (Nakahara et al., 1999). However, to follow the disappearance of these states in hot and dense matter in detail and locate the *melting temperature* for various hadronic excitations requires considerably more computing resources than are available today. Prior to the disappearance of a state, interactions with the thermal medium will lead to temperature and density dependent shifts of the resonance peaks, as well as a broadening of these peaks. To resolve the structure of spectral functions to such a degree that shifts in resonance peaks and broadening of the spectral curve become statistically significant, accurate numerical results for hadron correlation functions are required. As in the case of calculations of transport coefficients, large lattices are needed to generate information on the correlation functions at many different time separations.

Calculations of transport coefficients will provide fundamental insight into the structure of hot and dense matter. It will allow scientists to quantify aspects of the extent to which the phenomenologically successful modeling of heavy ion collisions has a solid foundation in QCD; i.e., whether a near-equilibrium QGP described by QCD indeed equilibrates rapidly and can be characterized as an almost perfect fluid. Detailed information on the spectral function would confirm whether or not the QGP is strongly coupled at RHIC, and by varying the temperature in the LQCD calculations, scientists will learn how much the temperature has to be increased before the plasma becomes weakly coupled. This question will be centrally important when comparing the heavy ion data obtained at the RHIC and the LHC experiments because the temperature in the latter experiment will be about a factor 1.5 to 2 higher.

These calculations will strongly influence the analysis of experimental data obtained in heavy ion collisions.

EQUILIBRATION CHALLENGE: FROM THE COLOR GLASS CONDENSATE TO THE QUARK-GLUON PLASMA

Determining the time required for a QGP to form after the onset of the collision (i.e., the thermalization time) and determining the physics processes that drive the QGP formation are among the most important outstanding problems in the area of ultrarelativistic heavy ion collisions. The success of ideal hydrodynamics in describing bulk observables—such as the elliptic flow of the matter created in noncentral collisions—implies the matter must have a short thermalization time compared to the overall time scale of the reaction. A proper understanding of the QCD dynamics that lead to rapid thermalization is important in determining the transport coefficients and the equation of state of hot QCD.

To describe the approach to equilibrium the following are required: 1) a firm understanding of the initial configuration of partons in the colliding nuclei and the process by which they are liberated from the nucleus at the onset of the collision, and 2) detailed models and simulations of the processes that take place during the early nonequilibrium phase of the collision leading to the formation of a thermalized QGP.

A large nucleus moving near the speed of light is a very dense system of gluons. It is believed that nonlinear effects in QCD lead to a saturation of the rapid growth of the gluon density in the colliding nuclei with beam energy and mass number *A* when the phase-space occupation number is (nonperturbatively) large, of order $1/\alpha_S$. The effective theory describing this nonlinear regime of QCD is the CGC (McLerran and Venugopalan1994a,1994b;, McLerran 1994b Kovchegov 1996). McLerran and Venugopalan (1994a) proposed an effective action incorporating high-gluon density effects that amount to solving the classical Yang-Mills equations where the large-momentum degrees of freedom in the nucleus act as sources of color charge for the small-momentum degrees of freedom.

The classical description of gluon saturation is modified at higher energies due to quantum loop corrections. To this end, a new set of equations, commonly referred to as the JIMWLK (Jalilian-Marian – Iancu – McLerran – Weigert – Leonidov - Kovner) equations (Jalilian-Marian 1997a,1997b; Jalilian-Marian 1997b, Iancu et al. 2001), have been derived from a Wilsonian Renormalization Group formalism. They describe the evolution of *n*-point functions in QCD with energy. The resulting equations are an infinite hierarchy of coupled differential equations that are difficult to solve analytically. Nevertheless, they can be written in a form which—in principle—allows them to be solved by lattice gauge-theory techniques (Rummukainen and Weigert 2004)).

The CGC has explained several theoretical and phenomenological aspects of high-energy interactions quite successfully. Nevertheless, many important properties of the CGC remain to be addressed quantitatively, and to be tested by comparing with experimental data from RHIC and other colliders.

In recent years is has been shown that early studies of the driving mechanism for the equilibration of quark-gluon matter in ultra-relativistic heavy ion collisions overlooked a crucial aspect of the dynamics of nonequilibrium plasmas, namely the possibility of plasma instabilities. Most importantly, it has been shown that these instabilities can produce plasma isotropization and thermalization on a time scale relevant to relativistic heavy ion collisions. The possibility of such non-Abelian plasma instabilities was first predicted in the mid-1990s (Mrowczynski 1993) by studying plasmas with an anisotropic momentum-space distribution. The resulting instability has been termed the chromo-Weibel instability. In recent years, this idea has received a significant amount of attention due to analytic and numerical advances. The first advance was to show instabilities predicted by Mrowczynski (1993) were generic and independent of the precise details of the assumed anisotropic distribution function (Romatschke and Strickland 2003). The second advance was the development of a gauge-invariant nonequilibrium effective action that included all non-Abelian interactions self-consistently (Mrowczynski et al. 2004). This second advance made it possible to numerically study the dynamics of the chromo-Weibel instability using established real-time lattice gauge theory techniques.

The goal of this PRD is to deliver the first real-time calculation of the collision of two heavy ions at high energy, retaining the complete three-dimensional structure of the fields of produced gluons (in impact parameter and rapidity space), energy-momentum conservation, and quantum evolution of the measure. The subsequent real-time evolution of the color fields following the initial impact will clarify the time scales and processes that lead to thermalization and formation of a QGP and the possible role played by non-Abelian gauge-field instabilities (plasma-turbulence). Distribution of the thermalized gluons in impact parameter and rapidity will provide much needed initial conditions for hydrodynamic modeling of the late stages of the collision, and could provide information on the equation of state and the viscosity of hot QCD matter. The project would also provide predictions for the effect of early-time nonequilibrium dynamics on important QGP observables such as jet quenching, anomalous transport, and fluctuations.